arXiv:1611.09408v2 [math.ST] 24 Dec 2016

A finite mixture model approach to regression under covariate misclassification

Michelle Xia[1] and P. Richard Hahn*[2]

[1]Division of Statistics, Northern Illinois University, Dekalb, IL, 60115, U.S.A.
[2]Booth School of Business, University of Chicago, Chicago, IL, 60637, U.S.A.

December 28, 2016

Abstract. This paper considers the problem of mismeasured categorical covariates in the context of regression modeling; if unaccounted for, such misclassification is known to result in misestimation of model parameters. Here, we exploit the fact that explicitly modeling covariate misclassification leads to a mixture representation. Assuming common parametric families for the mixture components, and assuming that the misclassification occurrence is independent of the response variable, the mixture representation permits model parameters to be identified even when misclassification probabilities are unknown. Previous approaches to covariate misclassification use multiple surrogate covariates and/or validation data on the magnitude of errors. Based on this mixture structure, we demonstrate that valid inference can be performed on all the parameters even when no such additional information is available. Using Bayesian inference, the method allows for learning from data combined with external information on the magnitude of errors when such information does become available. The method is applied to adjust for misclassification on self-reported cocaine use in the Longitudinal Studies of HIV-Associated Lung Infections and Complications (Lung HIV). We find a substantial and statistically significant effect of cocaine use on pulmonary complications measured by the relative area of emphysema, whereas a regression that does not adjust for misclassification yields a much smaller estimate.

Key words: Bayesian inference, Covariate misclassification, Generalized linear models, Markov chain Monte Carlo, Measurement error modeling, Mixture regression models.

1 Introduction

Misclassification refers to measurement error on categorical or binary variables, so that some observations within a data set might record a category other than the truth. For example,

*P. Richard Hahn, University of Chicago Booth School of Business, 5807 S. Woodlawn Avenue, Chicago, IL 60637, U.S.A. Email: richard.hahn@chicagobooth.edu

our motivating example arises because survey respondents are often reluctant to truthfully report drug use. When covariates or exposure variables are subject to misclassification, naive (unadjusted) estimation of regression effect can be biased (see, e.g., Liu and Liang [1991], Beavers and Stamey [2012], Wang and Gustafson [2014]).

Previous methods that adjust regression estimates for covariate misclassification use information from either multiple surrogates of the misclassified variable (see, e.g., Liu and Liang [1991]) or validation data on the misclassification probabilities (see, e.g., Chu, Gustafson, and Le [2010]). Recently, Shieh [2009] and Hubbard et al. [2016] studied a mixture representation for regression with misclassified categorical covariates, respectively for normal and binary response variables. Based on the mixture representation, the papers proposed frequentist methods adjusting for the misclassification, for situations when the reclassification probabilities are known (e.g., from previous studies).

In this paper, we expand the methods of Shieh [2009] and Hubbard et al. [2016] to study the general structure of mixture regression models resulting from covariate misclassification. The mixture representation demonstrates that model parameters can be identified even when none of the aforementioned side information is available. Specifically, the proposed method does not rely on additional information on the magnitude of errors, such as the reclassification probabilities assumed in Shieh [2009] and Hubbard et al. [2016]. Using Bayesian inference, external information on the reclassification probabilities can be incorporated through priors, when such information does become available. This helps us benefit from the capacity of existing methods in incorporating external information, while allowing additional statistical learning from data. Our approach can be applied to a wide class of regression models, in addition to the normal and binary cases studied in earlier papers, including generalized linear models and parametric survival models, provided the mixture components are identifiable. Our model allows for additional (correctly measured) covariates, provided that the misclassification is *nondifferential*, meaning that misclassification is conditionally independent of the response variable. We study the effectiveness of the mixture representation approach via an asymptotic analysis of efficiency and via a Monte Carlo study. The upshot of these investigations is that, if the effect size is large enough, our approach gives up little statistical efficiency relative to the case where the classification probabilities are known *a prior*. When the effect size is small, we demonstrate that the proposed model works as good as alternatives in cases where accurate information is available on the magnitude of errors. When the information is inaccurate, our method is able to learn this from the data, thereby reducing bias in the estimation.

Our motivating empirical example is the multi-center Lung HIV study (Crothers et al. [2011]) that was conducted during 2007 to 2013 by the National Heart, Lung, and Blood Institute (NHLBI). The goal of this study was to understand the relationship between human immunodeficiency virus (HIV) infection and pulmonary diseases. Data from the Lung HIV study and its sub-studies have been analyzed in the medical literature (Drummond et al. [2015], Leader et al. [2016]) in an effort to identify risk factors for pulmonary complications related to HIV infection. Other papers (e.g., Sigel et al. [2014], Depp et al. [2016]) assess HIV infection itself as a risk factor for pulmonary condition; Drummond et al. [2013], Lambert et al. [2015] focus on injection drug users.

For the Lung HIV study, self-reported traits were collected on use of illegal drugs such as crystal methamphetamine, crack cocaine, marijuana and heroin. Previous literature indicated a strong clinical connection between pulmonary diseases and use of recreational drugs, particu-

larly cocaine use (Yakel and Eisenberg [1995], Alnas, Altayeh, and Zaman [2010]; Fiorelli et al. [2016]). On the relationship between drug use and pulmonary function specifically in HIV-infected individuals, Simonetti et al. [2014] reported no significant effect of recreational drug use on pulmonary function, based on two sub-studies of the Lung HIV data (Multicenter AIDS Cohort Study and Women's Interagency HIV Study). The goal of our analysis is to revisit this question using the full Lung HIV study, while accounting for the covariate misclassification due to the probable inaccuracy of self-reported drug use. Due to the anticipated large effect size of cocaine use and the unavailability of external information, our method becomes a natural choice in adjusting for misclassification when assessing the cocaine effect.

The rest of the paper is organized as follows. In Section 2, we provide a general framework for regression models adjusting for covariate misclassification. In Section 3, we obtain and evaluate the asymptotic results on the efficiency loss for the normal model. In Section 4, we present a simulation study on finite samples. In Section 5, the methodology is applied to adjustment for misclassification in the self-reported cocaine use, when assessing its effect on lung density measures. Section 6 concludes the paper.

2 A finite mixture model representation

Throughout the paper, we use upper-case Roman letters to denote random variables such as Y, with the corresponding lower-case letter y denoting an observed value of the variable. We use lower-case bold letters such as $\mathbf{x}$ to denote a vector of random variables, and the notation $\boldsymbol{x}$ for an observed value of the random vector. Similarly, the bold version of parameters such as $\boldsymbol{\alpha}$ represents a vector of parameters. All vectors are column vectors, unless specified otherwise. Upper-case bold letters such as $\mathbf{P}$ are used to denote a matrix. For notational simplicity, we present a population model based on a single set of random variables, without involving an observed sample of size $n > 1$.

2.1 A finite mixture regression representation

Let V be a categorical variable taking $K \geq 2$ categories, which we will denote $(0, 1, \cdots, K-1)$, and let $P(V = k) = \pi_k \geq 0$ with $\sum_k \pi_k = 1$ denote the corresponding category probabilities. Denote by V^* the observed version of V with misclassification. The *misclassification probability* is defined as $p_{kj} = \mathrm{P}(V^* = j \,|\, V = k)$, with $\sum_j p_{kj} = 1$ for $k = 0, 1, \cdots, K-1$. Hence, the severity of misclassification can be represented with the classification matrix $\mathbf{P} = (p_{kj})$. Interested readers may refer to Buonaccorsi [2010], Gustafson and Greenland [2015], Yi [2016] for comprehensive reviews on the issue and adjustment of misclassification in categorical variables.

Here, we formulate the misclassification problem in the context of parametric regression, where the distribution of the response variable Y is modeled conditional on a categorical covariate V and additional (accurately measured) covariates denoted using a vector $\mathbf{x}$. We write the probability function of $(Y \,|\, V, \mathbf{x})$ as $f_Y(y \,|\, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\varphi}, V, \mathbf{x})$ for parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\varphi})$. Specifically, in the case of generalized linear models (GLMs) with a binary V, we have $g(\mathrm{E}(Y \,|\, V, \mathbf{x})) = \alpha_0 + \alpha_1 V + \mathbf{x}'\boldsymbol{\beta}$, with $\boldsymbol{\alpha} = (\alpha_0, \alpha_1)'$, a vector $\boldsymbol{\beta}$ containing the regression coefficients for $\mathbf{x}$, $\boldsymbol{\varphi}$ containing possible nuisance parameters such as a dispersion parameter, and $g(\cdot)$ being the link function.

Assume that the misclassification is nondifferential on Y, meaning that the occurrence of misclassification is conditionally independent of Y. Given the observed covariates V^* and $\mathbf{x}$, the conditional distribution of Y has mixture representation

$$f_Y(y \,|\, V^*,\, \mathbf{x}) = \sum_{j=0}^{K-1} q_j(V^*,\, \mathbf{x})\, f_Y(y \,|\, \boldsymbol{\alpha},\, \boldsymbol{\beta},\, \boldsymbol{\varphi},\, V = j,\, \mathbf{x}), \tag{1}$$

where the mixture weights are the *reclassification probabilities*, defined as $q_{kj}(V^*, \mathbf{x}) = \mathrm{P}(V = j \,|\, V^* = k,\, \mathbf{x})$; $k = 0,\, 1,\, \cdots,\, K-1$. The specific form of the $q_j(\cdot)$, which is a function from the domain of $(V^*,\, \mathbf{x})$ to $[0, 1]$, depends on the joint distribution of $(V,\, V^*,\, \mathbf{x})$. The K-by-K *reclassification matrix* is denoted $\mathbf{Q} = (q_{kj})$, with $\sum_j q_{kj} = 1$ for $k = 0,\, 1,\, \cdots,\, K-1$, where dependence on $(V^*,\, \mathbf{x})$ is left implicit. Equation (1) gives a weighted sum of K regression models, each corresponding to a particular unobserved value of V, with corresponding weights depending on the observed value of V^* (which can take the same K distinct values as V). This extends the mixture distribution representation in Shieh [2009] to a more general *mixture regression model* framework (see, e.g., Grün and Leisch [2008a]).

2.2 Identification

The model in (1) is a special case of a more general class of models known as mixture regression models with concomitant variables (Grün and Leisch [2008a]) or mixture-of-experts models (Jacobs et al. [1991]). This more general representation takes the form

$$f_Y(y \,|\, \mathbf{w},\, \mathbf{x}) = \sum_{j=0}^{K-1} \phi_j(\mathbf{w})\, f_Y(y \,|\, \boldsymbol{\omega}_j,\, \mathbf{w},\, \mathbf{x}), \tag{2}$$

where each $\boldsymbol{\omega}_j$ is a vector of component-specific parameters, and $\mathbf{w}$ and $\mathbf{x}$ are vectors of observed covariates. According to the first nomenclature, covariates $\mathbf{w}$ appearing in the mixture weights $\phi_j(\cdot)$ are called concomitant variables; according to the second, each regression component is referred to as an "expert" and the $\phi_j(\cdot)$ are referred to as gating functions. Although it is possible that the gating functions take more general forms, it is common to use a multinomial logit form

$$\phi_j(\mathbf{w}) = \frac{\exp\left(\nu_j + \mathbf{w}'\boldsymbol{\gamma}_j\right)}{\sum_{h=0}^{K-1} \exp\left(\nu_h + \mathbf{w}'\boldsymbol{\gamma}_h\right)}, \quad j = 0,\, 1,\, \cdots,\, K-1. \tag{3}$$

Identification of parameters in these models is a nontrivial issue, which is discussed in Hennig [2000], Grün and Leisch [2008b], and Jiang and Tanner [1999]. In these papers it is shown that models of the form in (2) are identifiable — meaning there is a unique mapping from probability functions $f_Y(y \,|\, V^*,\, \mathbf{x})$ to parameters $\boldsymbol{\Omega} = (\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_{K-1})$ — provided that several criteria are satisfied. We discuss these conditions next, explaining why the models studied in this paper satisfy them.

First, the family of component densities $f_Y(\cdot)$ must be identifiable. That is, given a finite mixture of distributions in this family, one must be able to uniquely decompose it into its

constituent components. This is known to be true of many common densities, including the normal, gamma, and Poisson distributions (Yakowitz and Spragins [1968], Atienza, Garcia-Heras, and Munoz-Pichardo [2006]; Atienza et al. [2007]) considered in this paper. Note that these results require that the mixture representation is *irreducible*, meaning that all K mixture components are distinct; this is true for our model if and only if $\alpha_1 \neq 0$.

Second, one must be able to order the parameter vectors $(\boldsymbol{\omega}_0 \prec \boldsymbol{\omega}_1 \prec \ldots, \boldsymbol{\omega}_{K-1})$ uniquely. This ordering serves to (arbitrarily) break the symmetry inherent to mixture representations, which is that they are invariant under permutations of the component labels. Jiang and Tanner [1999] provides a recipe for establishing such an ordering for exponential family models. Moreover, this condition is easily satisfied for (1) because only the (scalar) $\alpha_1 V$ term varies by component.

Regarding this second condition, it is important to note that in the misclassified covariate context the parameter ordering is not arbitrary. For example, in the case of a binary variable V, one component of the mixture model corresponds to the mean response when $V = 0$ and the other to the mean response when $V = 1$; clearly these designations have applied meaning and are not interchangeable. Fortunately, the relative ordering of these components will be preserved as long as the misclassification is not *systematic*, meaning that the probability of misclassification is higher than that of correct classification and that the direction (sign) of the slope parameter α_1 is known. See Weinberg, Umbach, and Greenland [1994] for additional discussion.

Third, one must parametrize the reclassification probabilities in terms of a *base category*, for example by setting $\nu_{K-1} = 0$ and $\boldsymbol{\gamma}_{K-1} = \mathbf{0}$. This breaks the invariance to location shifts of the gating parameters $(\nu_j, \boldsymbol{\gamma}_j')$.

Finally, Hennig [2000] points out that a mixture of regression models have additional criteria that must be satisfied for a given set of covariate values. There and in Grün and Leisch [2008b], concrete examples are produced which show that different parameters $\boldsymbol{\Omega} \neq \boldsymbol{\Omega}'$ can give the same likelihood evaluation if "[component] labels are fixed in one covariate point according to some ordering constraint, [but] labels switch in other covariate points for different parameterizations of the model." Note that Jiang and Tanner [1999] rule out this possibility by stipulating that $\mathbf{w} = \mathbf{x}$ defines a set with a non-null interior; by contrast, Grün and Leisch [2008b] provide a counter-example with a binary covariate. We satisfy this condition trivially — for any covariate values — because the misclassified covariate finite mixture in (1) defines mixture components with *parallel* hyperplanes, which share all parameters except the $\alpha_1 V$ term. Accordingly, component ordering is preserved across all covariate values.

The identifiability of (1) ensures that parameter estimates can be obtained without side information on the magnitude of misclassification, such as validation data or known reclassification probabilities required by the existing methods. Instead, the specific form of the model being used allows estimation even though V is not directly observed. Because this result may be counterintuitive, it is helpful to consider the normal response case with a dichotomous covariate and no concomitant variables. In that case, observed bimodality in the distribution of the response (for a fixed V^*) must be due to misclassification. Thus, simple visual inspection can be used to identify the misclassified observations, correct them, and proceed with the regression analysis. This simple process becomes much more difficult with concomitant variables, but a similar logic may be implemented through a formal likelihood-based analysis, which is the proposal of this paper.

3 Asymptotic evaluation

Based on the mixture representation in Equation (1), we can use the regular asymptotic theory for studying the efficiency loss due to the existence of covariate misclassification.

3.1 Fisher information

Without loss of generality, we study the situation where there is no additional covariates $\mathbf{x}$ in the model. In this case, the likelihood for a single observation of $(Y,\, V^*)$ can be written as

$$L^*(\boldsymbol{\alpha},\, \boldsymbol{\varphi},\, \mathbf{Q},\, \boldsymbol{\pi}^*) \propto \left[\prod_{k=0}^{K-1} (\pi_k^*)^{I_{(V^*=k)}}\right] \left[\sum_{j=0}^{K-1} q_{V^*j}\, f_Y(Y \mid \boldsymbol{\alpha},\, \boldsymbol{\varphi},\, V=j)\right], \tag{4}$$

where $\boldsymbol{\pi}^* = \mathbf{P}'\boldsymbol{\pi}$ is the category probabilities for the observed variable V^*, $I_{(A)}$ is the indicator function on whether the event A happens, and the reclassification probability only depends on V^* and j.

Denote by $\boldsymbol{\theta}$ the vector that contains all the parameters in $\boldsymbol{\alpha}$, $\boldsymbol{\varphi}$, $\boldsymbol{\pi}^*$ and $\mathbf{Q}$. The corresponding expected Fisher information matrix has the form $\mathcal{I}^*(\boldsymbol{\theta}) = E\left[(\partial l^*(\boldsymbol{\theta})/\partial\boldsymbol{\theta})^{\otimes 2} \,\middle|\, \boldsymbol{\theta}\right]$, where $l^*(\boldsymbol{\theta})$ is the log-likelihood function for a single observation.

For the asymptotic variance, we consider three *successively harder* scenarios for estimating the parameters in $\boldsymbol{\theta}_A$, the block of parameters of interest including $\boldsymbol{\alpha}$ and $\boldsymbol{\varphi}$.

The first case is *when perfect observations of* $(Y,\, V)$ *are available* (i.e., when there is no misclassification). Denote by $l(\boldsymbol{\alpha},\, \boldsymbol{\varphi}) = \log f_Y(Y \mid \boldsymbol{\alpha},\, \boldsymbol{\varphi},\, V)$ the log-likelihood function for a single observation of $(Y,\, V)$, and by $\mathcal{I}(\boldsymbol{\theta}_A)$ the corresponding expected Fisher information. Then the maximum likelihood estimator (MLE) $n\,\widehat{\boldsymbol{\theta}}_A \to MVN(n\,\boldsymbol{\theta}_A,\, Acov_1)$, as $n \to \infty$. The asymptotic variance covariance matrix can be written as

$$Acov_0 = [\mathcal{I}(\boldsymbol{\theta}_A)]^{-1}. \tag{5}$$

The next easiest scenario is when the reclassification matrix $\mathbf{Q}$ is *known*, along with the observed values of $(Y,\, V^*)$. Because in practice the reclassification matrix is typically unknown, this represents an optimistic scenario; the methods of Shieh [2009] and Buonaccorsi [2010] can be used in this case. For the MLE $\widehat{\boldsymbol{\theta}}_A$, the asymptotic variance covariance matrix is the corresponding block of covariance matrix in $[\mathcal{I}^*_{CC}(\boldsymbol{\theta})]^{-1}$. That is,

$$Acov_1 = \left\{[\mathcal{I}^*_{CC}(\boldsymbol{\theta})]^{-1}\right\}_{AA}, \tag{6}$$

where the parameter block $\boldsymbol{\theta}_C$ consists of the parameters in $\boldsymbol{\theta}_A$ ($\boldsymbol{\alpha}$ and $\boldsymbol{\varphi}$) and $\boldsymbol{\theta}_B$ ($\boldsymbol{\pi}^*$), and $\left\{[\mathcal{I}^*_{CC}(\boldsymbol{\theta})]^{-1}\right\}_{AA} = [I^*_{AA}(\boldsymbol{\theta})]^{-1} - I^*_{AB}(\boldsymbol{\theta})[I^*_{BB}(\boldsymbol{\theta})]^{-1} I^*_{BA}(\boldsymbol{\theta})$.

Finally, the most challenging scenario is when we observe only $(Y,\, V^*)$, and the reclassification matrix $\mathbf{Q}$ is *unknown*. In this case, previous methods cannot be used but our new mixture approach can be used. When the mixture distribution is identifiable (see, e.g., Yakowitz and Spragins [1968], Atienza et al. [2007]), the MLE $\widehat{\boldsymbol{\theta}}_A$ will have consistency and asymptotic normality, with the asymptotic covariance matrix given by

$$Acov_2 = \left\{[\mathcal{I}^*(\boldsymbol{\theta})]^{-1}\right\}_{AA}. \tag{7}$$

Given a parametric form for $f_Y(Y \mid \boldsymbol{\alpha}, \boldsymbol{\varphi}, V)$, the expected Fisher information can be evaluated numerically allowing to study the efficiency loss for inference on the parameters of interest, most importantly the regression coefficients corresponding to effects of interest.

3.2 Efficiency loss

Here we illustrate the efficiency loss for a normal linear model when there is one binary covariate V that is subject to nondifferential misclassification. Specifically, let $(Y \mid V) \sim N(\mu_V, \sigma^2)$, with the conditional mean $\mu_V = \alpha_0 + \alpha_1 V$, and constant conditional variance, σ^2.

Denote by $\widehat{\alpha}_1$ the MLE of the covariate effect α_1, and $\boldsymbol{\theta} = (\alpha_0, \alpha_1, \sigma^2, \boldsymbol{\pi}^*, \mathbf{Q})$. For the easiest scenario, when (Y, V) are observed, standard asymptotic theory gives $n \operatorname{Var}(\widehat{\alpha}_1) \to Avar_0(\boldsymbol{\theta})$, with a closed-form $Avar_0(\boldsymbol{\theta}) = \sigma^2/[\pi_1(1-\pi_1)]$. In order to obtain the forms of the asymptotic variance terms in Equations (6) and (7), we can write the log-likelihood for a single sample of (Y, V^*) as

$$\begin{aligned} l^*(\boldsymbol{\theta}) = & \log\left[q_{V^*0}\, f(Y \mid \alpha_0, \sigma^2) + (1 - q_{V^*0})\, f(Y \mid \alpha_0 + \alpha_1, \sigma^2)\right] \\ & + V^* \log(\pi_1^*) + (1 - V^*) \log(1 - \pi_1^*) + C, \end{aligned} \tag{8}$$

where C is a constant, and $f(Y \mid \alpha_0, \sigma^2)$ denotes the normal density function with mean α_0 and variance σ^2.

Based on the log-likelihood function, we can obtain the expected Fisher information matrix as

$$I^*(\boldsymbol{\theta}) = \pi_1^* E[s^{\otimes 2}(\boldsymbol{\theta}) | V^* = 1] + (1 - \pi_1^*) E[s^{\otimes 2}(\boldsymbol{\theta}) | V^* = 0],$$

where the score vector is given by

$$s(\boldsymbol{\theta}) = \left(\frac{\partial\, l^*(\boldsymbol{\theta})}{\partial\, \alpha_0}, \frac{\partial\, l^*(\boldsymbol{\theta})}{\partial\, \alpha_1}, \frac{\partial\, l^*(\boldsymbol{\theta})}{\partial\, \sigma}, \frac{\partial\, l^*(\boldsymbol{\theta})}{\partial\, \pi_1^*}, \frac{\partial\, l^*(\boldsymbol{\theta})}{\partial\, q_{00}}, \frac{l^*(\boldsymbol{\theta})}{\partial\, q_{10}} \right)', \tag{9}$$

with the analytical forms of the elements given in Section A in the supplement (Xia and Hahn [2016]).

Using numerical integration methods to evaluate each of the elements in $E[s^{\otimes 2}(\boldsymbol{\theta}) | V^* = v^*]$, it is possible to compute the asymptotic variance terms in Equations (6) and (7). Numerical integration is more stable after converting the infinite integral to a finite integral by transforming the variables using their cumulative distribution functions. We use the R function *integral*() from the **pracma** package, and found the adaptive methods based on the GaussKronrod quadrature (Davis and Rabinowitz [2007]) to be fast and reliable.

3.3 Numerical evaluation

Here, we study the *ratio of the asymptotic standard deviations*, denoted $Rasd_j = \sqrt{Avar_j / Avar_0}$ $(j = 1, 2)$, respectively for the two cases when the reclassification probabilities q_{01} and q_{10} are known and unknown. For the normal model, the ratio of the asymptotic standard deviations $Rasd_j$ only depends on three factors: the effect size α_1/σ, the severity of misclassification, and the binomial proportion π_1. In Figure 1 and Web Figures 1 through 2 in the supplement, we investigate five cases: $\alpha_1 = 0.3$, 0.5, 1, 2 and 5 with fixed $\sigma = 1$. In Figure 1, the binomial

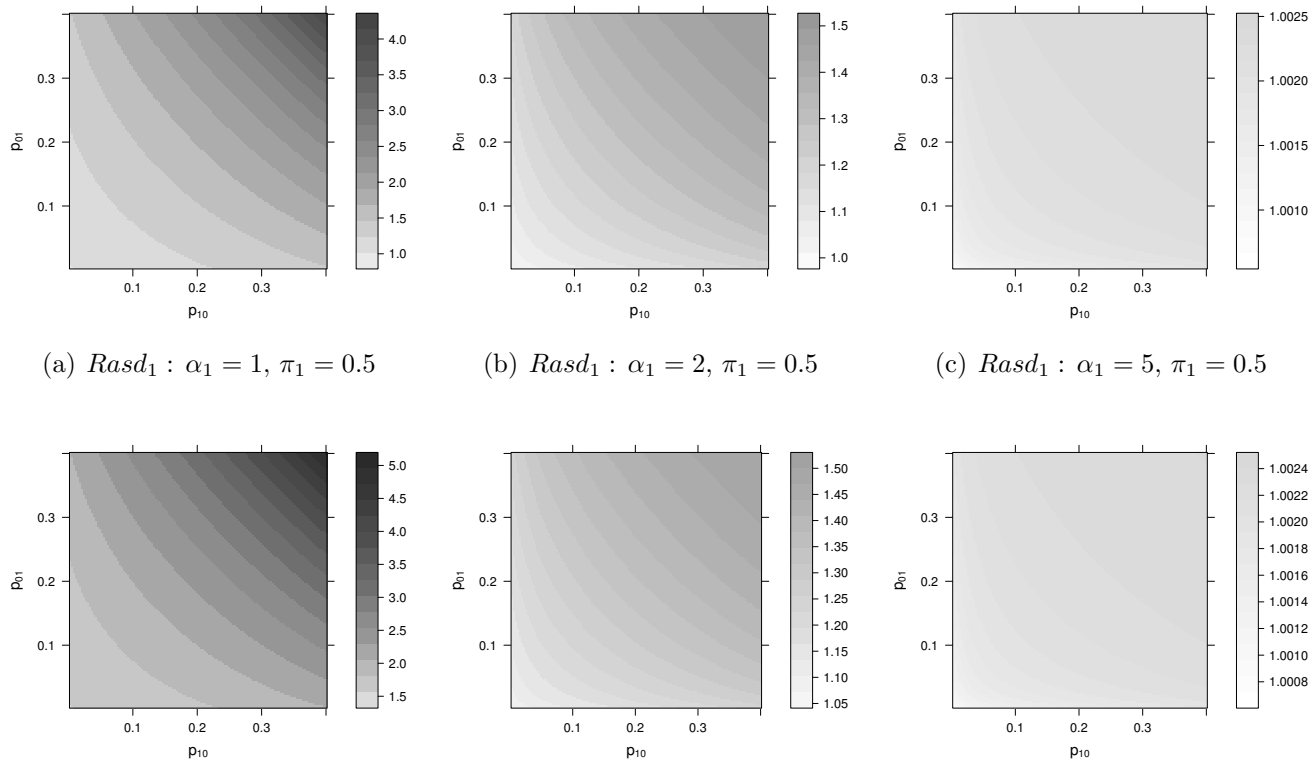

(a) $Rasd_1 : \alpha_1 = 1, \pi_1 = 0.5$ (b) $Rasd_1 : \alpha_1 = 2, \pi_1 = 0.5$ (c) $Rasd_1 : \alpha_1 = 5, \pi_1 = 0.5$

(d) $Rasd_2 : \alpha_1 = 1, \pi_1 = 0.5$ (e) $Rasd_2 : \alpha_1 = 2, \pi_1 = 0.5$ (f) $Rasd_2 : \alpha_1 = 5, \pi_1 = 0.5$

Figure 1: Efficiency loss for the effect α_1 with varying misclassification probabilities, under the normal model. The top panels present $Rasd_1$, the ratio of the asymptotic standard deviations for the case when the reclassification matrix is known versus that where the true status V is being observed. The bottom panels correspond to $Rasd_2$, that between the proposed model without knowing the reclassification matrix and the true model with accurately measured V.

proportion $\pi_1 = 0.5$ (Web Figure 2 in the supplement presents the results for $\pi_1 = 0.2$, Xia and Hahn [2016]).

Figure 1 shows that the efficiency loss is dominated by the effect size and the severity of the misclassification. For most of the scenarios, the ranges of the efficiency loss across different combinations of misclassification probabilities are similar, for the two cases when the reclassification matrix is known and unknown. Knowing the reclassification matrix offers more benefit in the cases when the effect size is small. When the effect size is large (i.e., ≥ 2), there is very little efficiency loss, regardless of whether the reclassification probabilities are known or not. A smaller or larger binomial proportion π_1 will result in mildly larger benefit from knowing the reclassification probabilities. In summary, large efficiency loss only occurs when the misclassification probabilities are large and the effect size is small. In these cases, the quality of the measurement is so questionable that a much larger confidence interval is needed, even for cases when we know the reclassification matrix. Note that on the boundaries of the figures, when one of the misclassification probabilities is 0 or 1, the ratio of the asymptotic standard deviations $Rasd_2$ will be infinite, owing to the nonidentifiability of mixture models in these scenarios (see, e.g., Li [2007], Xia and Gustafson [2016]).

4 Simulation studies

In order to illustrate the performance of the method in finite samples, we perform Monte Carlo simulation studies. Inference is performed using a fully Bayesian approach, based on Markov chain Monte Carlo (MCMC) samples from the posterior. This approach provides convenient statistical inference and the capability of incorporating prior information on the reclassification probabilities. For frequentist estimation, the `R` package **FlexMix** (Grün and Leisch [2008a]) can be used to mixture models using the EM algorithm, including the specific mixture models given in Equations (1) and (2).

Here, to illustrate our method, we consider two cases: the normal response case and the Poisson response case. More comprehensive studies for the normal, Poisson and also gamma models are given in the supplement (Xia and Hahn [2016]).

4.1 Normal model

4.1.1 Impact from effect size and external information

Here, we investigate the impact of the effect size and the (in)accuracy of external information concerning the reclassification probabilities on posterior estimates. In the normal linear model, we illustrate the convergence of the credible interval when there is an ordinal covariate V that is subject to nondifferential misclassification. In particular, we specify the conditional distribution of the response variable as $(Y \mid V) \sim N(\mu_V, \sigma^2)$, with the conditional mean having a linear form $\mu_V = \alpha_0 + \alpha_1 V$.

We assume that the ordinal covariate has a multinomial distribution, with probabilities $\boldsymbol{\pi} = (\pi_0, \pi_1, \pi_2) = (0.2, 0.3, 0.5)$ corresponding to the three values 0, 1 and 2. The following classification and corresponding reclassification matrices are assumed for obtaining the misclas-

sified sample of V^*.

$$\mathbf{P} = \begin{bmatrix} 0.80 & 0.15 & 0.05 \\ 0.10 & 0.70 & 0.20 \\ 0.05 & 0.15 & 0.80 \end{bmatrix} \quad \text{or} \quad \mathbf{Q} = \begin{bmatrix} 0.74 & 0.14 & 0.12 \\ 0.10 & 0.67 & 0.24 \\ 0.02 & 0.13 & 0.85 \end{bmatrix}.$$

After generating the sample of V from its multinomial distribution, we generate Y with variance $\sigma^2 = 4$, and regression coefficients (α_0, α_1) being either $(12, 2)$, $(12, 4)$ or $(12, 10)$. These three sets of parameters correspond to the same effect sizes as those in Figure 1 (i.e., α_1/σ, of 1, 2 and 5, respectively).

We illustrate the finite sample performance of four models with the sample sizes of 100, 400, 1,600, 6,400 and 25,600. The *naive model* refers to linear regression using the observed values of V^* as if there were no misclassification. The *true model* refers to linear regression using the correct classification V. The third model assumes that the reclassification matrix $\mathbf{Q}$ is known. Finally, our finite mixture model estimates the reclassification matrix based on the mixture representation.

For all four models, a normal prior with mean 0 and variance 100 is used for the intercept, and a gamma prior with parameters 0.0001 and 0.1 is used for the slope. The gamma prior represents our knowledge on the sign of the effect obtained from the naive estimate (i.e., the assumption that misclassification will not change the direction of the effect, in order to ensure the identifiability of the mixture model). For the normal model with an ordinal covariate, we place vague Dirichlet priors on the probability parameters in $\boldsymbol{\pi}^*$and each column of $\mathbf{Q}$, with concentration parameters of 1. The prior mean and standard deviation of the Dirichlet prior are 1/3 and 0.24. The above Dirichlet prior represents a situation when we have no external information on these probabilities. For the reclassification probabilities in the cases with $\alpha_1/\sigma =$ 1 and 5, we consider an additional set of informative priors for the proposed model. We use $60 \times q_{kj}$ as the concentration parameter for each of the reclassification probabilities, in order to match the prior mean with the true value. The resulting priors have small standard deviation ranging from 0.02 to 0.06. For the case where $\alpha_1/\sigma = 2$, we study the model performance when the external information the reclassification probabilities is inaccurate. Instead of knowing the true values, we assume the magnitudes of misclassification are under-estimated. In particular, we assume the estimates from external data are

$$\widehat{\mathbf{Q}} = \begin{bmatrix} 0.84 & 0.04 & 0.12 \\ 0.10 & 0.77 & 0.14 \\ 0.02 & 0.03 & 0.95 \end{bmatrix}.$$

Thus, we use the above estimates in the third model assuming known reclassification probabilities. For the proposed model, we consider an additional set of Dirichlet priors with the corresponding concentration parameters given by $60 \times \widehat{\mathbf{Q}}$ (i.e., representing an inaccurate guess). We take every 10th sample for all the models, after dropping the first 15,000 samples as burn-in. For a posterior sample of 5,000, the effective sample size is over 4,500 for all the models. The 95% equal-tailed credible intervals of the regression effect and other parameters from the four models are provided in Figure 2, and Web Figures 3 through 10 in the supplement.

Figure 2 reveals several intuitive patterns. First, the third panel shows that when the effect size is large, observing the true covariate status or knowing the reclassification probabilities

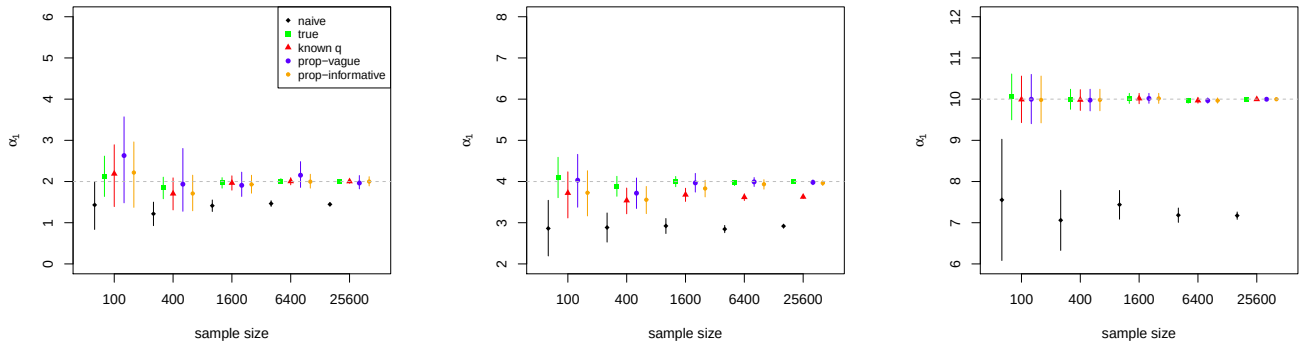

(a) $\alpha_1 = 2$ with accurate $\widehat{\mathbf{Q}}$ (b) $\alpha_1 = 4$ with biased $\widehat{\mathbf{Q}}$ (c) $\alpha_1 = 10$ with accurate $\widehat{\mathbf{Q}}$

Figure 2: Equal-tailed credible intervals of α_1 for normal responses with an ordinal V. With $\sigma^2 = 4$, the values of α_1 in the three panels correspond to an effect size of 1, 2 and 5. In the first panel, observe that the informative prior (gold) gives a posterior interval almost identical to the fixed $\widehat{\mathbf{Q}}$ case (red), both of which have notably smaller variance than the vague prior case, especially at lower sample sizes. In the second panel, observe that the informative prior is able to "unlearn" its prior bias when $\widehat{\mathbf{Q}}$ is incorrect, thus outperforming the fixed $\widehat{\mathbf{Q}}$ method (red) for sample sizes larger than 400. The final panel shows that when the effect size is large enough, all methods that account for misclassification are comparable.

confers no benefit for parameter inference — the Bayesian mixture analysis under a vague prior gives equivalent results. Second, the first panel shows, that informative priors give performance very similar to method which assumes known reclassification probabilities in terms of point estimate and posterior variance, and, by contrast, the vague prior has much wider posterior intervals; this shows that prior information can be incorporated successfully via a prior. Finally, the second panel shows that when the reclassification probabilities are incorrectly specified, an informative prior (centered at the incorrect value) can eventually overcome this bias, outperforming the method which fixes the reclassification probabilities at the wrong value; incorporating information about the reclassification probabilities via a prior is therefore a more conservative approach.

As expected, there is large bias in the naive estimates (which ignore misclassification). For the normal model, the variance can be dramatically over-estimated when misclassification is ignored. In the supplement (Xia and Hahn [2016]), it is shown that the reclassification probabilities can also be learned from the observed response and covariate, without further information. The learning is more efficient in the case when the effect size is large.

4.1.2 Impact from tail heaviness

For the normal model, we perform a sensitivity analysis on the impact from model misspecification. In particular, we consider scenarios where the response variables are generated from a Student-t distribution with the location and scale parameters denoted by μ_V and σ. We assume the same normal model as the previous subsection, and consider the large-effect case

with $\alpha_0 = 12$ and $\alpha_1 = 10$. The category and reclassification probabilities also take the same values.

We consider three cases in an attempt to understand the impact of tail heaviness, as well as that of larger variability caused by either a larger scale parameter or a smaller number of degrees of freedom for the Student-t distribution. For the first scenario, we assume that the number of degrees of freedom $\nu = 20$. We choose the scale parameter to be $\sigma = \sqrt{3.6}$, so that the variance is the same (4) as the large-effect case studied in the previous subsection. For the second and third scenarios, we increase the variance of the Student-t distribution by increasing the scale parameter σ to $\sqrt{6.48}$ or decreasing the number of degrees of freedom ν to 4. Both cases correspond to a variance of 7.2.

We consider five sample sizes, as in the previous subsection. We keep the MCMC and prior settings the same as the earlier case. The 95% equal-tailed credible intervals of the regression effect and other parameters are provided in Figures 3, and Web Figures 11 through 16 in the supplement.

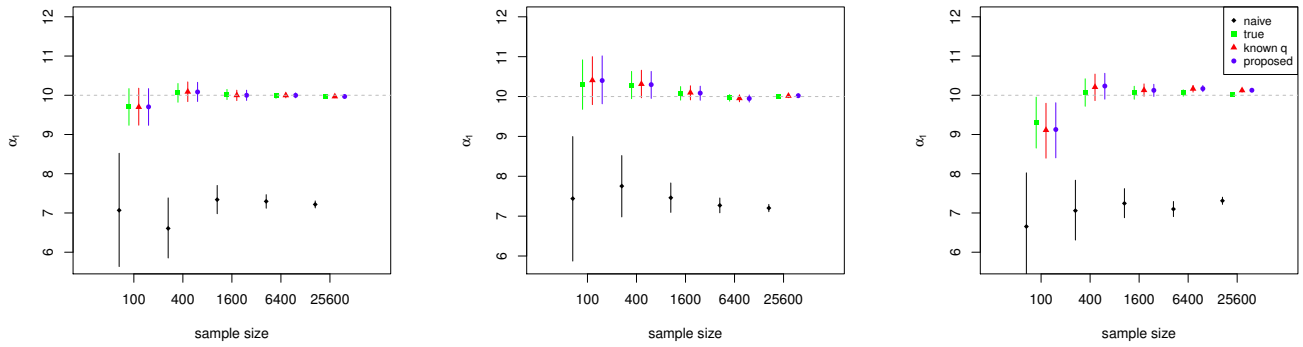

(a) $\sigma^2 = 3.6$, $\nu = 20$ (b) $\sigma^2 = 6.48$, $\nu = 20$ (c) $\sigma^2 = 3.6$, $\nu = 4$

Figure 3: Equal-tailed credible intervals of α_1 for Student t responses with an ordinal V. The variances for the three cases are 4, 7.2 and 7.2. With a slope $\alpha_1 = 10$, these sets of values correspond to an effect size of 5, 3.7 and 3.7, respectively.

We observe that: (1) mild tail heaviness does not seem to have an impact on the performance of the normal model; (2) increasing the scale parameter alone seems to lead to larger variability in the estimation (i.e., owing to a smaller effect size); (3) increasing the tail heaviness results in both small biases in the estimation and larger variability. The impact seems to be the same, regardless of whether the reclassification probabilities are known or not. The bias in the estimation is probably caused by the fact that lower (upper) tail values from both mixture components are more likely to be attributed to the group with a lower (higher) mean, when the tails are very heavy.

4.2 Poisson model

4.2.1 Impact from severity of misclassification

For the Poisson model, we study the impact of the severity of misclassification on finite sample inference. We assume there is a binary covariate V that is subject to nondifferential misclassification. We generate the sample of V using a Bernoulli trial with the probability $\pi_1 = 0.5$. In particular, we specify the conditional distribution of the response variable as $(Y \mid V) \sim Poisson\,(\mu_V)$, with the conditional mean $\mu_V = \exp(\alpha_0 + \alpha_1 V)$. The corresponding sample of Poisson counts Y is then generated from that of V, with the regression coefficients (α_0, α_1) assumed to be $(1.2, 1)$. Two different scenarios, $(p_{01}, p_{10}) = (0.1, 0.125)$ and $(p_{01}, p_{10}) = (0.125, 0.25)$, are assumed for the misclassification probabilities for obtaining the corresponding sample of V^*.

Similar to the case of the normal model, we explore the finite sample performance at the sample sizes of 100, 400, 1,600, 6,400 and 25,600. We compare the equal-tailed credible intervals from the naive model, the true model, the model with the known reclassification matrix $\mathbf{Q}$, and the proposed model as defined in the earlier subsection. For all models, independent normal priors with mean 0 and variance 10 are used for the regression coefficients α_0 and α_1. For the parameters p_{01}, p_{10} and π_0, independent uniform priors on $(0, 1)$ are used. The MCMC details are the same as the normal case. The 95% equal-tailed credible intervals of the regression effect and other parameters from the four models are provided in Figure 4, and Web Figures 17 and 19 in the supplement.

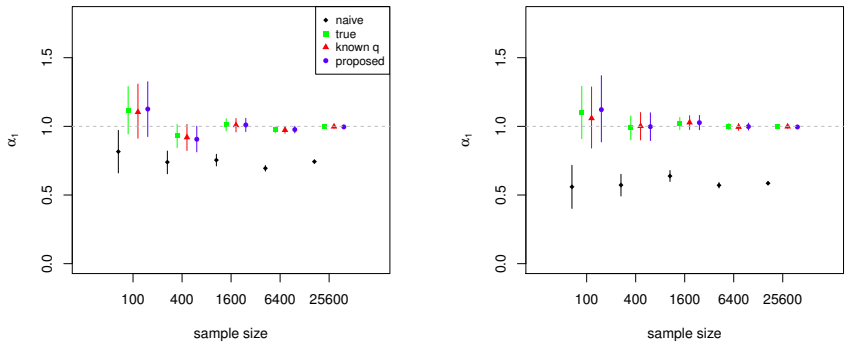

(a) $p_{10} = 0.125$ (b) $p_{10} = 0.25$

Figure 4: Equal-tailed credible intervals of α_1 for Poisson responses with a binary V.

The figures show that the naive model demonstrates larger attenuation bias in the second scenario, when the misclassification probabilities are larger. Because the effect size is large, there is little efficiency loss regardless of the severity of the misclassification and knowing the reclassification matrix offers little extra efficiency. For both cases, the credible intervals of the misclassification probabilities, p_{01} and p_{10}, are seen to converge to the true values with increasing sample size. Knowing the reclassification probabilities helps with the efficiency of estimating their misclassification counterparts. These results confirm the conclusions from an asymptotic evaluation based on the observed Fisher information.

4.2.2 Impact from zero inflation

For count data, it will be interesting to perform a sensitivity analysis on the impact from model misspecification when there is zero inflation. Here, we assume that the response variable Y is generated from a zero-inflated Poisson (ZIP) distribution. We assume that the conditional mean of the ZIP distribution is given by $\mu'_V = (1-w)\mu_V = \exp(\alpha_0+\alpha_1 V)$, with w being the percentage of additional zeros and μ_V being the mean of the Poisson component. For the sensitivity analysis, we take the less-severe misclassification scenario with $(p_{01}, p_{10}) = (0.1, 0.125)$. We study two scenarios on zero inflation with $w = 5\%$ and 10%, a fixed percentage of zeros across the $V = 0$ and $V = 1$ groups. We consider two cases with $(\alpha_0,\ \alpha_1) = (1.2,\ 1)$ and $(0,\ 1)$, with the conditional distributions of Y presented in Web Figure 20.

Again, we study the performance of the four models under five sample sizes. The prior and MCMC settings are the same as those in the previous subsection. The 95% equal-tailed credible intervals of the regression effect and other parameters are provided in Figures 5, and Web Figures 21 through 24 in the supplement.

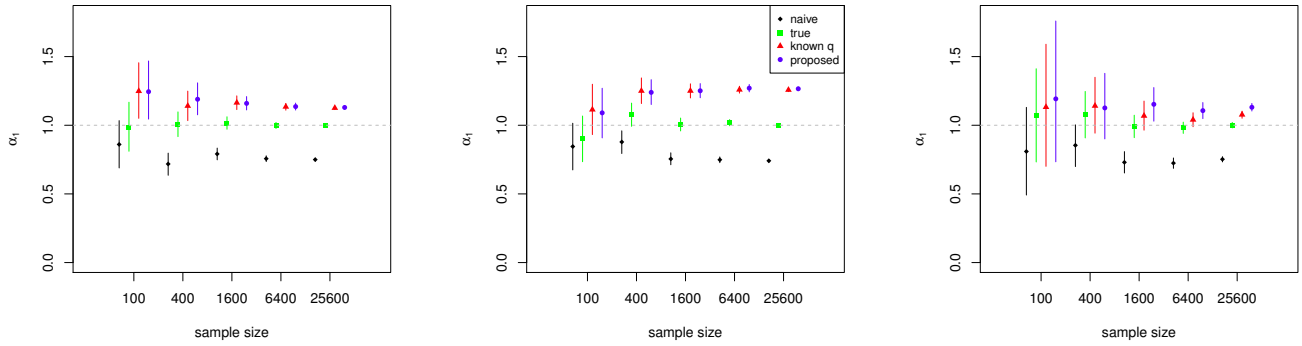

(a) $\alpha_0 = 1.2$, $w = 5\%$ (b) $\alpha_0 = 1.2$, $w = 10\%$ (c) $\alpha_0 = 0$, $w = 5\%$

Figure 5: Equal-tailed credible intervals of α_1 for ZIP responses with a binary V. The first panel corresponds to a variance to mean ratio of 1.19 for the $V = 0$ case and 1.47 for the $V = 1$ case. The second panel corresponds to a variance to mean ratio of 1.36 for the $V = 0$ case and 2.01 for the $V = 1$ case. The last panel corresponds to a smaller effect size, with a variance to mean ratio of 1.06 for the $V = 0$ case and 1.15 for the $V = 1$ case.

Figure 5 shows that the effect is over-estimated for the mixture models regardless whether the reclassification probabilities are known or not. It seems the two methods treat some or all of the additional zeros as coming from the $V = 0$ group that has a smaller mean. The main observations are: (1) the bias increases with the severity of zero inflation; (2) knowing the reclassification offers more benefit in the case with a smaller effect size; (3) the reallocation of the additional zeros seems to have led to a small decrease in the variability of estimation, when comparing to the results in the earlier subsection.

5 Empirical analysis: Lung HIV study

5.1 Data

The Longitudinal Studies of HIV-Associated Lung Infections and Complications (Lung HIV, Crothers et al. [2011]) was a collaborative multi-site study conducted between 2007 and 2013 by the National Heart, Lung, and Blood Institute (NHLBI). The study undertook data and specimen collection from eight HIV and pulmonary studies associated with NHLBI. The study's intent was to advance knowledge on HIV-related pulmonary diseases, motivated by the high incidence of serious pulmonary complications in HIV patients.

The Lung HIV study included adults over the age of 18 who were diagnosed of HIV and self-reported smoking. There were a total of 904 participants included in the study who had a computed tomography (CT) performed at the baseline. In particular, the CT provided lung density measures including the relative area of emphysema (RA) below -910 Hounsfield units (HU, RA-910) and the 15th percentile density in HU (PD15) that have been regarded as effective assessments of the extent and severity of chronic obstructive pulmonary disease (COPD, see, e.g., Soejima et al. [2000], Shaker et al. [2011]). Self-reported traits on use of recreational drugs such as cocaine and heroin were collected on a voluntary basis. Use of illegal drugs, particularly cocaine, is a known risk factor for pulmonary disease in the general population (see, e.g., Yakel and Eisenberg [1995], Alnas et al. [2010], Simonetti et al. [2014], Fiorelli et al. [2016]). Baseline demographic information including age, gender, race, education, work status, family size and living arrangement are also included.

After excluding records with missing values in the corresponding responses and covariates, the sample size ranges from 436 to 442 for the six distinct response variables. The numbers of participants who reported a positive status are 400 (male), 346 (exposure to smoking), 173 (white), 263 (cocaine use) for the binary traits. The age of the participants ranges from 21 to 75, and the number of cigarettes each participant smokes daily ranges from 0 to 45. The education variable is ordinal with 6 categories, with the value increasing with the level of education. We perform a logarithms transformation for the RA-910 and RA -600 to -250 variables, and all the variables seem to satisfy the normality assumption after the transformation.

5.2 Related literature

Several recent papers have analyzed data from the Lung HIV study or its sub-studies, focused on understanding pulmonary complications related to HIV infection. Among them, Drummond et al. [2015] and Leader et al. [2016] appear to be the only two that analyzed the multi-center Lung HIV data. Both papers study risk factors such as age and CD4 cell count for lung complications in HIV-infected individuals. Additional papers use data from sub-studies of the broader Lung HIV study. Among them, papers such as Sigel et al. [2014] and Depp et al. [2016] study the association between HIV infection and pulmonary diffusing capacity and lung density measures, using data from the Examination of HIV-Associated Lung Emphysema and Veterans Aging Cohort Study that included both HIV-infected and non-infected participants. Drummond et al. [2013] and Lambert et al. [2015] assess the association between HIV infection and lung function decline among injection drug users, using the AIDS Linked to the IntraVenous Experience cohort data on injection drug users.

Of these previous papers, the most similar to our analysis is Simonetti et al. [2014], which examines the association between pulmonary function and drug use in HIV-infected individuals, using data from the Multicenter AIDS Cohort Study and Women's Interagency HIV Study cohorts. They found no significant effect of recreational drug use on pulmonary function. There are two factors that may have limited the study's ability to detect a significant drug effect. First, the study had a total sample of only 184 patients, with 84 drug users. This sample size is small for the binary logistic regression model that was used. Second, intentional misrepresentation on drug use status was unaccounted for, although likely, due to legal/social desirability concerns.

Our approach addresses both of these concerns. By using the multi-center Lung HIV data we have a larger sample size, and our approach naturally handles the possibility of misclassification of the drug use variable.

5.3 Analysis and results

We are interested in assessing the effect of cocaine use in a continuous regression model on measures of lung density, after accounting for other risk factors and potential misclassification of the cocaine use variable. The response variables of interest include lung density, PD15, mean density, RA below -910, RA below -856, and RA between -600 to -250 HU. Among the demographic and risk factors available in the study, we include age, gender, ethnicity, exposure to smoking (ExpSmk), the number of cigarettes smoked each day (NoCigs), education (Edu), and use of crack cocaine (CraCoc) as covariates.

We perform an adjusted analysis using the proposed approach, and an unadjusted analysis using regular Bayesian normal model. We implement both the unadjusted and adjusted models in WinBUGS, through the R package **R2WinBUGS**. For all models, we run three chains with thinning of 6 and a burn-in of 15,000. Normal priors with mean 0 and variance 100 are used for the regression coefficients, except for the intercepts of the PD15 model, for which we set the prior variance to 10,000 due to its large scale. For the normal precision variable, we specify a vague gamma prior with parameters 1.5 and 2. For the reclassification probabilities, q_{01} and q_{10}, we assume they do not depend on the other covariates. We use a beta prior with parameters 1 and 9, corresponding to a prior mean and standard deviation of 0.1. The prior puts more weights at the values near zero, so we are assuming the chance of misclassification is small, unless data suggest otherwise. As suggested by Ferrari et al. [2008], evidence from more informative studies may be incorporated through priors in a Bayesian model when such information is available. For all parameters, we take a posterior sample of size 5,000 with an effective size over 4,500.

The prior and posterior densities for the reclassification probabilities are presented in Web Figure 25. the posterior mean and standard deviation for the reclassification probabilities (q_{01}, q_{10}) are [0.67(0.05), 0.22(0.04)] for the RA-910 model. The posterior standard deviation is much smaller than the prior standard deviation, indicating significant learning in the reclassification probabilities. The RA-910 model indicates a much larger probability of false negatives than that of false positives. This is expected due to the social undesirability and legal concern surrounding recreational drug use. The posterior mean of q_{01} from the PD15 model is slightly larger (0.11) than the prior mean (0.1), indicating very weak learning of the reclassification probability when the effect size is small. For the density model, the posterior mean and stan-

dard deviation of the reclassification probabilities are the same as those of the priors. This is presumably due to the nonidentifiability of the model when the effect size of the cocaine use variable is zero.

The 95% equal-tailed credible intervals for the covariate effects are given in Figure 6 and Web Figure 26. For the two variables of RA−910 and RA −600 to −250, we fit the models on the logarithm scale, so we present the exponential of the slope that represents the multiplicative effect on the median of these variables. For the other variables, we present the regression slope that can be interpreted as the additive effect on the mean of the variables.

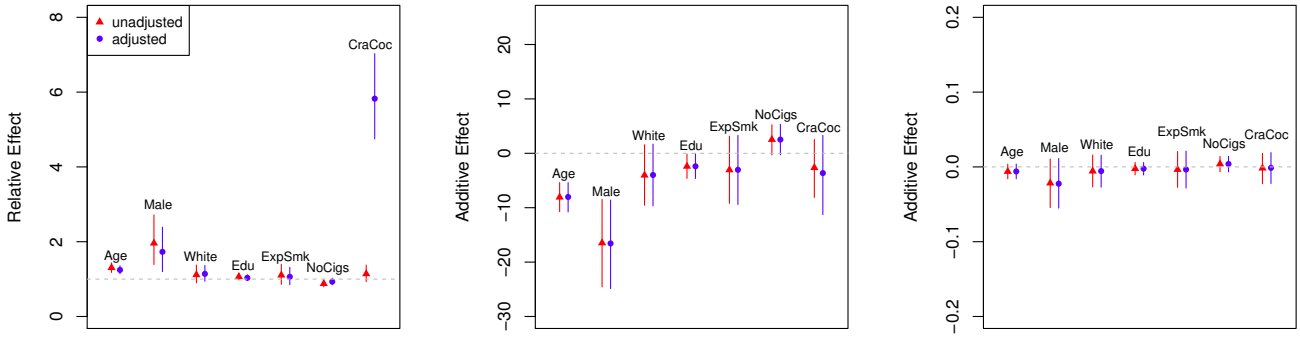

(a) RA−910 (b) PD15 (c) Density

Figure 6: Equal-tailed credible intervals of covariate effects on the lung density measures including the RA−910, PD15, and lung density. A log transformation is performed on RA−910 in order to satisfy the normality assumption. Hence, the exponential of the coefficient corresponds to the relative effect on the median of RA−910. The age effect corresponds to every 10-year increase in the age, and the cigarette effect corresponds to every 10-cigarette increase per day.

In Figure 6, we observe that the directions of the effect for all covariates match those found in the previous literature. Of particular note is that the number of cigarettes a participant smoked has an estimated effect that is opposite to that of the age and cocaine use variables. This is consistent with the findings from Shaker et al. [2011]: inflammation from smoking may mask the presence of emphysema on CT. After adjusting for misclassification in the self-reported cocaine use status, we observe a significant effect of cocaine use on worsening lung complications measure by the RA below −910 HU. The estimated effect size for the model is around 2.6, confirming it is a case where the proposed model is efficient. Since the relative effect is the exponential of the regression coefficient, the adjustment results in a large impact on the estimated relative effect. For the other two response variables, PD15 and lung density, the adjustment does not result in a significant difference in the estimated effects.

In summary, our analysis suggests that the high incidence of cocaine use in HIV-infected individuals is probably a major contributing factor for the severe or fatal pulmonary complications that have been observed in the target population.

5.4 Sensitivity analysis

Here, we perform a sensitivity analysis on how the results concerning RA−910 can be affected when the reclassification probabilities are permitted to vary as a function of other covariates.

Specifically, we model the reclassification probabilities q_{01} and q_{10} using the logit model given in (3). In the binary logit model, we include covariates race, education and age, which are plausibly associated with cocaine use or misclassification status. We investigate six logit models including either one, two, or all three covariates.

MCMC and prior settings are similar to those in the previous subsection. In particular, for the regression coefficients of the logit models, we select the priors in an effort to match the unconditional prior mean and standard deviation for reclassification probabilities, those of the $beta(1, 9)$ prior we assumed in the previous subsection. For the single-covariate logit models, we standardize the corresponding covariates before using a $N(-2.5, 0.04)$ prior for the intercept and a standard normal prior for the slope. For the multi-covariate logit models, the priors are selected in a similar manner. Figure 7 shows the 95% equal-tailed credible intervals of the covariate effects on RA−910 for the models we have considered. The prior and posterior distributions of the regression coefficients for the logit models are presented in Web Figure 27.

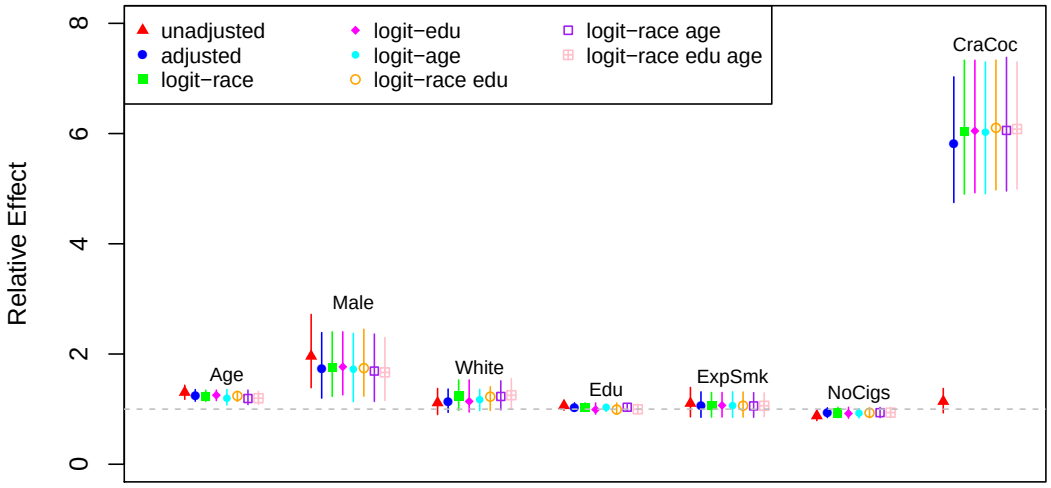

Figure 7: Equal-tailed credible intervals of covariate effects on RA−910, for the sensitivity analysis with logit models on reclassification probabilities q_{01} and q_{10}.

Figure 7 shows that regression on reclassification probabilities leads to a mild increase in the estimated cocaine effect, with the cost of a mild increase or decrease in those of the corresponding covariates. From Web Figure 27, none of the covariate effects seems to differ significantly from zero. Overall, the basic findings appear to be robust to the modeling assumptions governing the reclassification probabilities.

6 Conclusions

In this paper, we formulate regression models with a misclassified categorical covariate as mixture regression models. The finite mixture representation enables us to perform valid statistical inference on all parameters in the model, including the regression coefficients and reclassification probabilities, without requiring extra sources of information on the misclassification probabilities or the true covariate value. By evaluating the efficiency loss caused by the misclassification, we showed that the efficiency loss is dominated by the effect size, the severity of misclassification, and the distribution for the categorical covariate of concern. For cases where the effect size is large, not observing the true covariate value and not knowing the reclassification probabilities contributes little to loss of efficiency. Therefore, conducting inference using the mixture represetnation allows for reliable inference on model parameters even when there is no information or validation data available on the misclassification or the true status. Further, furnishing side information on the reclassification probabilities via an informative prior protects against the possibility that the side information is wrong. Applying the methodology to the Lung HIV study data, we confirmed that the adjustment of misclassification in the self-reported cocaine use results in an estimated effect that is substantially larger than that indicated by the unadjusted analysis. Our analysis indicates a significant effect of crack cocaine use in worsening lung complications measured by the relative area of emphysema, after adjusting for other known risk factors.

Acknowledgments

This manuscript was prepared using Lung HIV Research Materials obtained from the NHLBI Biologic Specimen and Data Repository Information Coordinating Center and does not necessarily reflect the opinions or views of the Lung HIV and the NHLBI.

References

M. Alnas, A. Altayeh, and M. Zaman. Clinical course and outcome of cocaine-induced pneumomediastinum. *The American Journal of The Medical Sciences*, 339(1):65–67, 2010.

N. Atienza, J. Garcia-Heras, and J. Munoz-Pichardo. A new condition for identifiability of finite mixture distributions. *Metrika*, 63(2):215–221, 2006.

N. Atienza, J. Garcia-Heras, J. Munoz-Pichardo, and R. Villa. On the consistency of MLE in finite mixture models of exponential families. *Journal of Statistical Planning and Inference*, 137(2):496–505, 2007.

D. P. Beavers and J. D. Stamey. Bayesian sample size determination for binary regression with a misclassified covariate and no gold standard. *Computational Statistics & Data Analysis*, 56 (8):2574–2582, 2012.

J. P. Buonaccorsi. *Measurement error: models, methods, and applications*. CRC Press, 2010.

R. Chu, P. Gustafson, and N. Le. Bayesian adjustment for exposure misclassification in case-control studies. *Statistics in Medicine*, 29:994–1003, 2010. doi: 10.1002/sim.3829.

K. Crothers, B. W. Thompson, K. Burkhardt, A. Morris, S. C. Flores, P. T. Diaz, R. E. Chaisson, G. D. Kirk, W. N. Rom, and L. Huang. HIV-associated lung infections and complications in the era of combination antiretroviral therapy. *Proceedings of the American Thoracic Society*, 8(3):275–281, 2011.

P. J. Davis and P. Rabinowitz. *Methods of numerical integration*. Courier Corporation, 2007.

T. B. Depp, K. A. McGinnis, K. Kraemer, K. M. Akgün, E. J. Edelman, D. A. Fiellin, A. A. Butt, S. Crystal, A. J. Gordon, M. Freiberg, et al. Risk factors associated with acute exacerbation of chronic obstructive pulmonary disease in HIV-infected and uninfected patients. *AIDS*, 30(3):455–463, 2016.

M. B. Drummond, C. A. Merlo, J. Astemborski, M. M. Marshall, A. Kisalu, J. F. Mcdyer, S. H. Mehta, R. H. Brown, R. A. Wise, et al. The effect of HIV infection on longitudinal lung function decline among injection drug users: a prospective cohort. *AIDS*, 27(8):1303–1311, 2013.

M. B. Drummond, L. Huang, P. T. Diaz, G. D. Kirk, E. C. Kleerup, A. Morris, W. Rom, M. D. Weiden, E. Zhao, B. Thompson, et al. Factors associated with abnormal spirometry among HIV-infected individuals. *AIDS*, 29(13):1691–1700, 2015.

P. Ferrari, R. J. Carroll, P. Gustafson, and E. Riboli. A Bayesian multilevel model for estimating the diet/disease relationship in a multicenter study with exposures measured with error: the EPIC study. *Statistics in medicine*, 27(29):6037–6054, 2008.

A. Fiorelli, M. Accardo, F. Rossi, and M. Santini. Spontaneous pneumothorax associated with talc pulmonary granulomatosis after cocaine inhalation. *General Thoracic and Cardiovascular Surgery*, 64(3):174–176, 2016.

B. Grün and F. Leisch. FlexMix version 2: finite mixtures with concomitant variables and varying and constant parameters. *Journal of Statistical Software*, 28:1–35, 2008a.

B. Grün and F. Leisch. Finite mixtures of generalized linear regression models. In *Recent advances in linear models and related areas*, pages 205–230. Springer, 2008b.

P. Gustafson and S. Greenland. Misclassification. *Handbook of Epidemiology, 2nd Ed.*, pages 639–658, 2015.

C. Hennig. Identifiability of models for clusterwise linear regression. *Journal of Classification*, 17(2):273–296, 2000.

R. A. Hubbard, E. Johnson, J. Chubak, K. J. Wernli, A. Kamineni, A. Bogart, and C. M. Rutter. Accounting for misclassification in electronic health records-derived exposures using generalized linear finite mixture models. *Health Services and Outcomes Research Methodology*, pages 1–12, 2016.

R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.

W. Jiang and M. A. Tanner. On the identifiability of mixtures-of-experts. *Neural Networks*, 12 (9):1253–1258, 1999.

A. A. Lambert, G. D. Kirk, J. Astemborski, S. H. Mehta, R. A. Wise, and M. B. Drummond. HIV infection is associated with increased risk for acute exacerbation of COPD. *Journal of Acquired Immune Deficiency Syndromes*, 69(1):68–74, 2015.

J. K. Leader, K. Crothers, L. Huang, M. A. King, A. Morris, B. W. Thompson, S. C. Flores, M. B. Drummond, W. N. Rom, and P. T. Diaz. Risk factors associated with quantitative evidence of lung emphysema and fibrosis in an HIV-infected cohort. *Journal of Acquired Immune Deficiency Syndromes*, 71(4):420–427, 2016.

P. Li. *Hypothesis testing in finite mixture models*. PhD Dissertation, University of Waterloo, 2007.

X. Liu and K. Y. Liang. Adjustment for non-differential misclassification error in the generalized linear model. *Statistics in Medicine*, 10(8):1197–1211, 1991.

S. B. Shaker, T. Stavngaard, L. C. Laursen, B. C. Stoel, and A. Dirksen. Rapid fall in lung density following smoking cessation in COPD. *COPD: Journal of Chronic Obstructive Pulmonary Disease*, 8(1):2–7, 2011.

M. S. Shieh. *Correction methods, approximate biases, and inference for misclassified data*. PhD Dissertation, University of Massachusetts at Amherst, 2009.

K. Sigel, J. Wisnivesky, S. Shahrir, S. Brown, A. Justice, J. Kim, M. Rodriguez-Barradas, K. Akgün, D. Rimland, G. S. Hoo, et al. Findings in asymptomatic HIV infected patients undergoing chest computed tomography testing: Implications for lung cancer screening. *AIDS*, 28(7):1007–1014, 2014.

J. A. Simonetti, M. R. Gingo, L. Kingsley, C. Kessinger, L. Lucht, G. Balasubramani, J. K. Leader, L. Huang, R. M. Greenblatt, J. Dermand, et al. Pulmonary function in HIV-infected recreational drug users in the era of anti-retroviral therapy. *Journal of AIDS & Clinical Research*, 5(11):1–14, 2014.

K. Soejima, K. Yamaguchi, E. Kohda, K. Takeshita, Y. Ito, H. Mastubara, T. Oguma, T. Inoue, Y. Okubo, K. Amakawa, et al. Longitudinal follow-up study of smoking-induced lung density changes by high-resolution computed tomography. *American Journal of Respiratory and Critical Care Medicine*, 161(4):1264–1273, 2000.

D. Wang and P. Gustafson. On the impact of misclassification in an ordinal exposure variable. *Epidemiologic Methods*, 3(1):97–106, 2014.

C. A. Weinberg, D. M. Umbach, and S. Greenland. When will nondifferential misclassification of an exposure preserve the direction of a trend? *American Journal of Epidemiology*, 140(6): 565–571, 1994.

M. Xia and P. Gustafson. Bayesian regression models adjusting for unidirectional covariate misclassification. *Canadian Journal of Statistics*, 44(2):198–218, 2016.

M. Xia and P. R. Hahn. Supplement to "A finite mixture model approach to regression under covariate misclassification". 2016.

D. L. Yakel and M. J. Eisenberg. Pulmonary artery hypertension in chronic intravenous cocaine users. *American Heart Journal*, 130(2):398–399, 1995.

S. J. Yakowitz and J. D. Spragins. On the identifiability of finite mixtures. *The Annals of Mathematical Statistics*, 39(1):209–214, 1968.

G. Y. Yi. *Statistical Analysis with Measurement Error or Misclassification: Strategy, Method and Application*. Springer, 2016.